Generation

11

Learning Objectives

After completing this reading, you should be able to

- Describe the operating characteristics of the following generation technologies: coal-fired; natural gas-fired, including combustion turbine and combined cycle gas turbine units; oil-fired and dual-fired generation; nuclear; hydropower; wind power; and solar, including PV and CSP
- Calculate and relate the heat rate and the thermal efficiency of a generating plant, and explain factors that can impact a plant's heat rate and its thermal efficiency
- Construct and interpret both a short-term and a long-term operating cost curve; explain factors that impact the operating cost of a generating plant
- Apply the heat rate and other cost factors to calculate the marginal cost of electricity for a generating plant
- Compare and calculate the capacity factor and availability factor, and explain how these factors vary for different types of generating plants

Excerpt is Chapter 3 of Introduction to Electricity Markets, *by John Parsons.*

Chapter 3

Generation

3.1 Introduction

This chapter introduces the main electricity generation technologies. It describes the key operating constraints and calls attention to the particular features of each technology that make it well suited to meeting certain characteristics of load and ill suited for others. The chapter also introduces you to a couple of key metrics on performance, cost and utilization. Finally, it introduces you to storage technologies.

3.2 Generation Technologies

Generating plants convert **primary energy** into electrical energy. The chemical energy in fossil fuels such as coal, natural gas and oil are among the main sources of primary energy. The nuclear energy in uranium is another primary energy source. The kinetic energy in wind and hydro resources is yet another, as is solar radiation. Figure 3.1 shows the shares of different primary energy sources used to generate electricity in the countries of the Organization for Economic Cooperation and Development (OECD) in 2014. Coal-fired generation accounted for 32% of the total. Natural gas-fired generation, which has been increasing its share in recent years, accounted for 24%. Oil-fired generation, which has been decreasing its share over many years, still accounted for 3%. Nuclear generation accounted for 18%. Renewables as a group, which include hydro-, wind- and solar power, accounted for 23%, and are now increasing their share significantly.

The figure also shows the shares in the non-OECD countries. Notably, coal has a larger share in the non-OECD, accounting for nearly 50% of total generation. Nuclear accounts for a smaller share in the non-OECD, but renewables account for a comparable share.

Each primary energy source has its distinguishing characteristics which create advantages and disadvantages:

- Coal is abundant across the globe, providing a plentiful supply of low cost fuel. It is easy to combust and use in generation, and it is easy to store at the generator. On the other hand, it is a very dirty fuel, contributing to a variety of pollution problems or requiring extra investment in mitigation. Emissions of particulates, sulfur dioxide (SO_2), nitrogen oxides (NO_x) and mercury are all important problems. CO_2 emissions are also a major issue. Coal combustion results in greater CO_2 emissions than oil and natural gas per unit of heat output because of its higher ratio of carbon to hydrogen and because the efficiency of a typical gas-fired generator is higher than the efficiency of a typical coal-fired generator.

- Natural gas is cleaner burning, and can be used in a variety of different types of generators, providing extra flexibility to an electricity system. However, historically, natural gas resources have had limited geographic availability. Because of its gaseous form and low energy density, natural gas was historically transported by pipeline. In recent decades, trade in liquefied natural gas (LNG) has expanded dramatically, which has significantly expanded its geographic availability. Either way, delivery costs

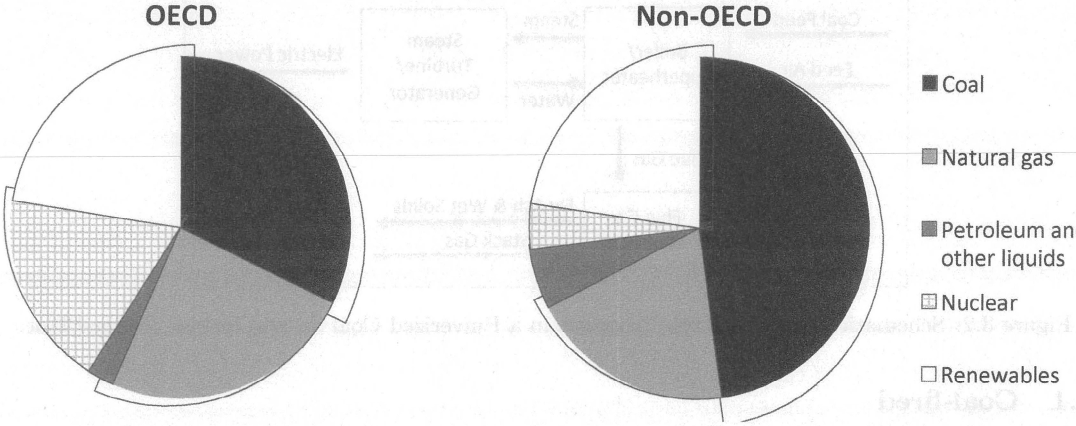

Figure 3.1: Fuel Shares in Electricity Generation in OECD and non-OECD countries, 2014.

Data source: International Energy Agency, World Energy Statistics 2016.

represent a large fraction of the total cost of natural gas.[1] It is also difficult to store in large volume except in a few locations with especially favorable geologies. For these and other reasons, the price of natural gas has proven very volatile. While burning natural gas produces less CO_2 emissions than coal, these emissions are still a significant contributor to atmospheric carbon and therefore to the danger of global warming.

- Oil's share in electricity generation is declining because products made from oil, such as gasoline or diesel, are uniquely valuable as a fuel for transportation in cars, trucks, aircraft and shipping. However, the ability to store oil more readily than natural gas has encouraged its use as a back-up fuel for electric generation to meet rare peak load demand or to substitute when natural gas supplies are locally limited.
- Nuclear generation is a complex technology that requires a large up-front capital investment. It brings with it important safety concerns, as well as concerns regarding disposal of the waste fuel rods. On the other hand, it is generally a very reliable source of electricity, and its fuel can be easily stockpiled so that there is no threat to supply. It also produces no CO_2 emissions, which is increasingly considered its main advantage.
- Existing renewable generation is primarily hydropower, although both wind and solar are increasing dramatically. Like nuclear, renewables require a relatively large upfront capital investment, but only a modest ongoing operating expense. Indeed, the fuel cost is zero. The availability of renewable resources varies greatly with geography. Certain regions have plentiful hydro or wind resources, while others have less. Although solar resources also vary by geography, a sufficient quantity is widely available to promise that with technological advancement solar generation could become a major source of electricity in almost all regions. Hydropower often provides a valuable source of stored energy, while wind and solar suffer from the intermittency of the resource and must be complemented with dispatchable generation, usually from fossil fuel-fired generators. Renewables produce no CO_2 emissions, although each type of renewable generation has specific environmental impacts that must be managed.[2]

---

[1] The Future of Natural Gas, MIT, 2011, Chapter 1.

[2] The reservoirs of some hydropower projects may produce methane emissions, which is an even more potent greenhouse gas than CO_2.

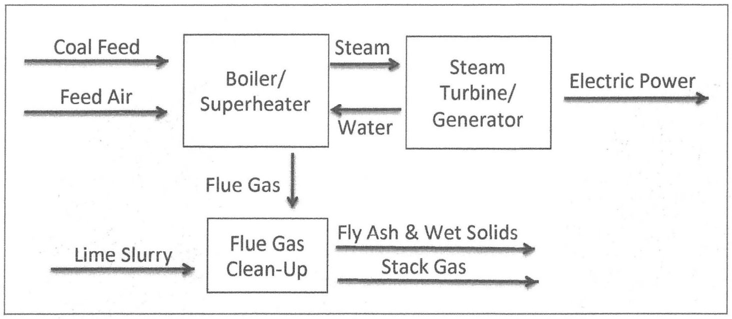

Figure 3.2: Schematic Diagram of the Processes in a Pulverized Coal Steam Turbine Power Plant.

3.2.1 Coal-fired

Steam Turbine plants

The primary coal-fired generating technology in widespread use is a **steam turbine** generator (ST) fueled by air-blown pulverized coal (PC). Figure 3.2 shows the three main components of a typical coal-fired generator: (i) coal combustion and steam generation, (ii) the steam turbine and electric generator, and (iii) ash and flue gas clean-up. The coal is ground to a fine powder and injected into the furnace with air where it is burned. Dry, saturated steam is generated in the boiler tubes and is heated further in the superheater. This high pressure, superheated steam drives a steam turbine coupled to an electric generator. The low-pressure steam exiting the steam turbine is condensed, and the condensate pumped back to the boiler for conversion into steam. The flue gas from the boiler passes through flue gas clean-up units to remove particulates, SOx, and NOx according to the local regulations. The stack gas exiting the clean-up section typically contains 10-15% CO_2.[3] Ash is removed from the furnace.

Steam turbine units are **dispatchable**–they can be turned on and off, and once on, their output in megawatts can be turned up and down. However, they also have important constraints on this dispatchability. They require a **minimum start time** before they can begin producing electricity. Once operating, they have not only a maximum, but also a **minimum generation level**, and they have a maximum **ramp rate**–the rate at which it can move from a lower level of generation to a higher level of generation, as well as a maximum rate at which they can ramp down. When the unit is brought down, it has a **minimum downtime** before it can be started back up again. Therefore, steam turbine units have proven to be good baseload plants or as mid-merit units following load through most of its predictable and gradual daily and seasonal cycles, but poor at matching unpredictable or quick ramps of load up and down. They are also good at providing certain types of security and stability to the system (often labeled ancillary services), which will be described in Chapter 4 on Transmission.

Coal is very easy to transport, and easy to stockpile at the power plant, which usually limits the danger that service could be interrupted. Coal is not a homogeneous fuel. The four main **ranks** in order of their heat content, from highest to lowest, are:

- Anthracite coal;
- Bituminous coal;
- Subbituminous coal; and
- Lignite coal.

The highest categories have heat content in the neighborhood of 24,000 kJ/kg and above, while the lowest is less than 17,000 kJ/kg.[4] Anthracite is used primarily in the metals industry. The other three grades are all used to generate electricity, among other things. Coals also vary along a number of dimensions, including moisture, ash, and sulfur contents. These characteristics impact how efficiently the power plant performs, as well as the cost of transportation and of emissions control. Despite the ease of transport, the cost of transport can become a significant share of the fuel cost and impact the choice of grade and other decisions. So, for example, lignite coals, with their low heat content and therefore relatively high transport costs, are usually only used by power plants in the immediate vicinity of where they are mined.

---

[3] The Future of Coal, MIT, 2007, Chapter 3.

[4] International Energy Agency, Energy Statistics Manual, 2005.

Coal-fired generation produces several types of hazardous wastes, including coal ash, and emissions to the air of nitrogen oxides (NO_x), sulfur, particulates, carbon monoxide (CO), carbon dioxide (CO_2), unburned hydrocarbons (UHC), volatile organic compounds (VOC), as well as of arsenic, mercury, lead and other toxic metals. Management of these wastes is usually subject to government regulation. Chapter 9 will discuss the economic management of some of these emissions in greater detail. Another environmental consideration is the amount of water required for cooling. Air cooling is feasible and used at plants located in arid climates, but it is more expensive.

Integrated Gasification Combined Cycle plants

An alternative coal-fired generation technology is integrated gasification combined cycle (IGCC). The coal is first transformed into a pressurized gas, known as syngas. The syngas is then processed, and the processed syngas is then burned to generate electricity. The pre-combustion processing of the syngas removes some of the elements like sulfur that in a PC unit are removed post-combustion. The electricity generation can be done in a combined-cycle plant, which has efficiency benefits that we discuss below under natural gas-fired generation. IGCC is a new, complex, and, as yet, little-used technology. However, advocates claim it promises a cheaper path for reducing the CO_2 emissions from coal-fired generation.

3.2.2 Natural Gas-fired

While Figure 3.2 shows a coal-fired steam turbine power plant, a steam turbine plant can just as well be fired by other fuels, including natural gas. However, historically coal has generally been a cheaper alternative for steam turbine systems. Natural gas can also be used to drive other types of power plants that offer their own advantages relative to a steam turbine system.

Combustion Turbines

In a **combustion turbine** (CT) or gas turbine power plant, the combustion of natural gas combined with air is used directly to drive the turbine, without the use of an intermediate steam cycle. A CT power plant is like a jet engine. It has three main sections: (i) the compressor, which draws air into the engine and pressurizes it, (ii) the combustion system, that injects fuel in combustion chambers where it mixes with air and is burned at high temperature, producing a stream of high temperature and high pressure gas, and (iii) the turbine, which is an array of stationary and rotating blades through which the gas travels and expands, spinning the rotating blades which drive the compressor to draw more air into the system and which also spin a generator to produce electricity.

A CT plant is usually less efficient than a steam turbine plant, and so CTs are seldom the most cost-effective way to provide baseload power. But, a CT plant has two important advantages. First, the capital cost is lower per unit of capacity, so it is a very cost effective provider of peaking power as we shall see in Chapter 5 on System Optimization. Second, CTs are very flexible and they can be used to quickly ramp power up or down as needed, so they can respond to some load requirements that steam turbine plants have difficulty meeting.

Combined-cycle Gas Turbine plants

A **combined-cycle gas turbine** (CCGT) plant gets its name from the fact that it combines the steam turbine plant and the combustion turbine plant into one. The two plants employ different thermodynamic cycles–the former is the Rankine Cycle, and the latter is the Brayton Cycle. Hence the name combined cycle. First, natural gas is used to fire a CT. Second, the waste heat from the turbine is captured and used to power a conventional steam turbine in a heat recovery steam generator (HRSG). The hot gases leaving the HRSG are treated to control CO and NO_x emissions and then discharged into the atmosphere. These plants are also called natural gas combined-cycle (NGCC) plants, or just combined-cycle plants.

CCGTs improve efficiency of the overall system by taking waste heat from the first plant and putting it to productive use in the second. The system can operate in different **configurations**: for start-up, the gas turbine runs alone, with the steam turbine disconnected using a hydraulic clutch. Single-shaft systems have one gas turbine, one steam turbine, one generator and one HRSG, with the gas turbine and steam turbine coupled to the single generator on a single shaft. Multi-shaft configurations have one or more gas turbine generators and HRSGs supplying steam through a common header to a separate single steam turbine generator.[5] More generally, some number of combustion turbines and steam turbines are configured to work together, and these can be made to operate in different configurations.

Each configuration has a determined generating region, operating efficiency, and ramping capability. There are operating constraints on switching from one configuration to another.

Fuel deliverability is a more critical issue for natural gas-fired plants than for many other types of generation, including coal. Many natural gas plants receive fuel from a pipeline system, and there is often little storage capacity available in the neighborhood of the power plant. Therefore, throughput capacity on the pipeline is critical to the availability of the generator. In many regions, gas pipelines also deliver natural gas for city and town heating systems, and these users may have priority when pipeline capacity is strained. Some natural gas plants are located near seaports and receive gas piped from nearby liquified natural gas (LNG) tanker terminals. These terminals have a limited supply of storage available, but rely on a flow of shipments aligned with the fuel usage at the plant and other users supplied by the terminal. Whether the natural gas is delivered by pipeline or ships, without careful management action, shortages of deliverable supply can arise that constrain power generation.

Natural gas is usually delivered in a standard grade, with relatively few impurities compared to coal. It burns more easily and cleanly. Nevertheless, it produces emissions of NO_x, CO, UHC, and CO_2 that may require management according to local regulations. The steam cycle in a combined-cycle plant usually requires a significant amount of water for cooling, although air cooling is possible.

3.2.3 Oil-fired and Dual-fired

Refined oil products such as heating oil can also be used as the fuel driving a steam turbine plant. In many regions and countries, oil has been a preferred alternative since natural gas has not been available. Oil products are easy to transport and are available globally, while natural gas is not. In most locations, storage for oil is also more readily available than storage for natural gas. However, where natural gas is available, it is often a cheaper alternative for power generation. Natural gas also burns more cleanly than oil, and this is another reason gas is sometimes preferred. Consequently, the share of oil-fired generation has declined in many countries. In part this is due to the fact that little new oil-fired capacity has been built as other capacity has expanded, and in part this is due to old oil-fired steam turbine plants being repurposed to burn natural gas.

Some formerly oil-fired steam turbine plants have been repurposed as **dual-fired** plants: they can burn either natural gas or oil. These plants burn natural gas the vast majority of time, and only burn oil when natural gas is unavailable for some reason. Preserving the option to burn oil provides added security for customers of the electricity system. These plants must maintain a stock of oil fuel readily available to the plant. Local environmental regulations may be crafted to forbid the use of oil most of the time, in order to obtain the environmental benefit of cleaner natural gas, but still allow the oil to be burned in a few emergency hours when the electricity is essential. Some exclusively oil-fired units can provide this sort of system security, too–they sit idle most of the time and only operate in emergency hours.

---

[5] http://electrical-engineering-portal.com/an-overview-of-combined-cycle-power-plant

Combustion turbine technology can be adapted to using diesel as a fuel instead of natural gas. This can provide the peaking, ramping and system security benefits of combustion turbines to a location where natural gas is unavailable.

Smaller diesel generators, which operate like an automobile engine, are often employed for systems that are not grid-connected or that have an unreliable connection, and as back-up generation for essential services in the rare event of a blackout or similar disconnection.

3.2.4 Nuclear

Nuclear power plants also adapt the basic design of a steam turbine system, but instead of burning coal, natural gas or oil in a furnace, they 'burn' uranium in a reactor. While we do speak of "burning" uranium, it is the fissioning of atoms that releases the energy. Heat from the reactor drives the steam/water cycle which drives the turbine.

Nuclear plants are very capital intensive, while the operating costs are usually very low. Therefore, nuclear plants are typically employed as baseload units. Since most nuclear plants are not intended to be load following, they are usually not designed to adjust their level of output more frequently. However, a nuclear plant can technically be designed to follow load, and in France, where nuclear plants have historically provided approximately 3/4 of the country's electricity, some nuclear plants do follow load.

Nuclear plants typically run for more than a year between refuelings, and the fuel is easily stockpiled, so there is little danger that lack of fuel will make the plant unavailable.

Nuclear plants are free of the major pollutants common to fossil fuel-fired generation, such as NO_x, CO, UHC, and CO_2. However, the safe management of radioactive materials is essential, especially to avoid the dangerous release of radioactive contaminants to the soil, atmosphere or water. Fuel waste is often stored on-site at very low cost and/or shipped either for geologic disposal or to be processed for recycling. The availability of sufficient water for cooling is an important consideration in siting plants. The safe operation of nuclear plants is a matter of public concern and debate.

3.2.5 Hydro

Hydropower is a simple and widely used generation technology that harnesses the kinetic energy of water flowing downhill. A dam or inlet concentrates water flow into a pipe leading to the turbines that translate the water's kinetic energy to an electrical generator. Many hydropower projects design the dam to create a large **reservoir** that stores water across seasons or even years, enabling the generator to operate near full capacity most of the time. In contrast, a **run-of-river** project is one that operates without a large reservoir, although it may have a small headpond.

Hydropower facilities differ in (i) **head height** (the vertical distance between the intake and turbine), (ii) discharge or **flow-rate**, and (iii) total capacity. The head height and flow-rate are important determinants of the optimal turbine design.[6] Hydropower facilities range widely in scale. The large scale dam systems are often among the largest power plants in the world. China's Three Gorges plant has a total capacity of 22.5 GW. Brazil's Itaipu plant, on the border with Paraguay, has a total capacity of 14 GW. At the other end of the spectrum, many run-of-river plants have a capacity of less than 1 MW.

The operating costs of hydropower plants are often quite low. For plants above 10 MW of capacity, the average cost is $45/kW/year, while for small plants in the range of 1-10 MW of capacity, the average cost is $45-50/kW/year, and for the smallest plants with a capacity below 1 MW, the average cost is $45-250/kW/year.[7] How these figures translate to a cost per MWh produced depends upon the capacity factor: a plant with a high capacity factor will allocate this annual cost over a larger number of units of generation as compared to a similar plant of the same size but with a lower capacity factor. Notice that the range of operating cost is very high. Like other metrics for hydropower plants, this one depends on the particular features of the facility and resource.

---

[6] International Energy Agency (IEA) & International Renewable Energy Agency (IRENA) Hydropower Technology Brief, 2015.

[7] International Energy Agency (IEA) & International Renewable Energy Agency (IRENA) Hydropower Technology Brief, 2015.

Reported operating costs of hydropower facilities focus on direct out-of-pocket costs needed to run the facility. Since they are largely fixed, they do not play a large role in how the facility is operated. The environmental costs associated with a facility are more difficult to assess and to reduce to a monetary value. However, it is often the environmental impacts that determine how a facility is operated. A hydropower facility may be mandated to allow a certain amount of water to flow at all times of the year. It may be mandated to allow the flow to vary in a pattern that respects important seasonal features of the natural system. It must be managed with an eye to the full annual patterns of rainfall and drought.

We will discuss the capital costs of hydropower facilities later, when we discuss capacity investments of all types of generation. However, it is worth noting here that the environmental and social impacts associated with how a region's water resources are managed are very significant so that the decision to construct a facility involves more than a choice about electricity generation. Even the choice about electricity generation is often more than just a question of the immediate cost. Since hydropower facilities are often a huge investment, they involve larger issues of regional economic development that must be examined together with the immediate cost of electricity.

3.2.6 Wind

The basic features of a wind turbine are straightforward. The blades intercept the wind and translate the wind's energy into rotation of a first shaft. A gearbox translates the slow rotation of the first shaft into the high speed rotation of a second shaft which drives an electric generator. Accompanying these basic elements are devices for monitoring wind speed and direction and regulating the orientation of the blades and the operation of the system, such as a pitch system, yaw drive, wind vane, ananometer and controller. All of this equipment is perched atop a tower, often made of tubular steel, with a secure foundation. The most widely used wind turbines for power generation have three blades on a horizontally oriented shaft, much like the blades of an aircraft engine. Other designs exist, too. Some come with a different count and placement of blades, and some have a vertically oriented shaft.

The capacity of an individual wind turbine can be anything up to 8 MW, although this maximum has been increasing with the development of technology and materials. The average size for grid-connected wind turbines is usually a little below 2 MW, although both the rotor sweep and the power rating of installed wind turbines has been increasing in recent years. Multiple wind turbines are often assembled carefully in an array known as a windfarm, which includes mid-voltage power lines to collect the power from each tower and deliver it to a substation where it is transformed to a high voltage for injection onto the grid. The towers are spaced to minimize each tower's interference in the airflow to other towers.

The quality of available wind resources varies dramatically geographically. Wind speeds of between 6 and 14 meters per second (13-31 miles per hour) are generally required. The wind resource is often measured in terms of the **specific power**, which is reported in units of W/m^2. Consistent wind speed is optimal. Turbulent wind impairs the performance of turbines, and therefore is not desirable. Most wind turbines are sited onshore. Wind conditions offshore are often better, but installation and maintenance of offshore wind turbines is much costlier. Different wind turbine designs are suitable for different wind conditions.

The potential generation of a wind farm is determined by the installed capacity and by the amount of the wind resource. A key challenge is the **intermittency** of the resource. This happens at different time scales. At very short frequencies, the output from individual turbines can be volatile around the mean, much as load shows volatility at short frequencies as seen in Figure 2.7. However, the scale of these fluctuations is much smaller than for load, and where many turbines are installed over a dispersed geographic territory, the effect is to smooth out the impact of the fluctuations.[8] In addition, it is possible to design the power electronics of a wind turbine to control these fluctuations to some degree.

At longer time scales–daily, seasonally, and longer–the problem is fundamentally different. The amount of specific power varies with the weather. On average, on-shore winds are strongest at night and weaker during the day, although not all nights have good winds. Winds can often be stronger in the winter than the summer, and occasional spells of especially hot, still weather can last for many days over wide territories. The intermittency of the resource translates into intermittency of the electricity generated from installed wind farms. It is feasible to **curtail** generation at a wind farm, which could be useful in periods when plentiful winds coincide with periods of low electricity demand. Given the small scale of capacity currently installed in most regions, this practice is uncommon. However, if the capacity installed increases significantly enough, it could become necessary on occasion. Of course, on the downside of intermittency, when there is a shortfall of wind resources it is not possible to increase production.

---

[8] This smoothing effect would not hold where there is a single large wind farm on a unique location.

On top of the intermittency, is the problem of unpredictability. While some of the weather related variations are predictable, many are difficult to predict, even for as short a horizon as a single day. Technology can be helpful here: significant improvements in the predictability of wind power have been made in recent years. Nevertheless, a certain amount of residual uncertainty is impossible to remove.

3.2.7 Solar

Solar Photovoltaic

The main solar power technology is **solar photovoltaics** (PV) which use solar cells to convert sunlight directly into electricity. Photovoltaics have been studied since the early 1800s. Famously, PV installations have been employed since the 1950s to power satellites and other vehicles in space. They are uniquely valuable in certain off-grid uses, where other forms of power generation are difficult to supply. Until the last 15 years, they had not been used very frequently for grid-connected generation. However, the cost of an installed PV system has fallen dramatically in recent decades, and PV systems are increasingly being used as a source of grid-scale carbon-free electricity. Once installed, they operate at ambient temperatures, without moving parts and with minimal maintenance.[9] They also produce no other local air pollutants.

A grid-connected PV installation consists of (i) PV modules, each containing many individual solar cells, (ii) balance-of-system (BOS) electronics hardware, such as combiner boxes, inverters, and transformers, and (iii) BOS installation hardware, such as racking, wiring, disconnects and enclosures. When the front surface of a solar cell is illuminated by sunlight, photons are absorbed and their energy is transferred to an electron and a positively charged counterpart called a hole. The internal electric field of the cell creates a DC current. The module gathers the current of the many cells, and the BOS electronic components convert this to an AC current which is delivered to the grid.

A module is typically 1 meter by 1.5 meters by 4 centimeters. It contains 60-96 individual solar cells which are 15 square centimeters (6 square inches), a backsheet for environmental protection, and an aluminum frame for mounting. Each cell is capable of producing 4-5 watts under peak illumination (W_p), and the typical peak power ratings for the full module range from 260 W to 320 W. Most modules are crystalline silicone based, although other materials are used currently and there is ongoing research into new materials and types of cells. A major objective is to improve the absorption of light, which would increase the production of electricity.[10]

A PV installation is an array of many modules. These can be placed on residential or commercial rooftops, and can be ground mounted at any scale. A coarse categorization is: (i) residential systems – up to 10 kilowatts (kW), (ii) commercial systems – ranging between 10 kW and 1 megawatt (MW), and (iii) utility scale systems – larger than 1 MW. Theoretically, PV modules operate with equal efficiency as their installation into arrays is scaled up: a 10-square-meter (m^2) array is no less efficient that a 10-square-kilometer (km^2) array. However, while the same modules operate with equal efficiency no matter the size of the array, the balance-of-system costs are often much higher per unit for smaller arrays. Installing 20 MW of modules on a single large plot of flat ground is less costly per unit than installing the same quantity on 200 commercial rooftops, which is in turn less costly per unit than installing the same quantity on 4,000

---

[9] Most installations are fixed-tilt arrays in which the physical orientation of the module is fixed. Some other installations adjust the tilt to track the sun and improve the amount of solar irradiation received by the module. Obviously, these installations do have moving parts.

[10] This section relies heavily on the MIT Future of Solar study.

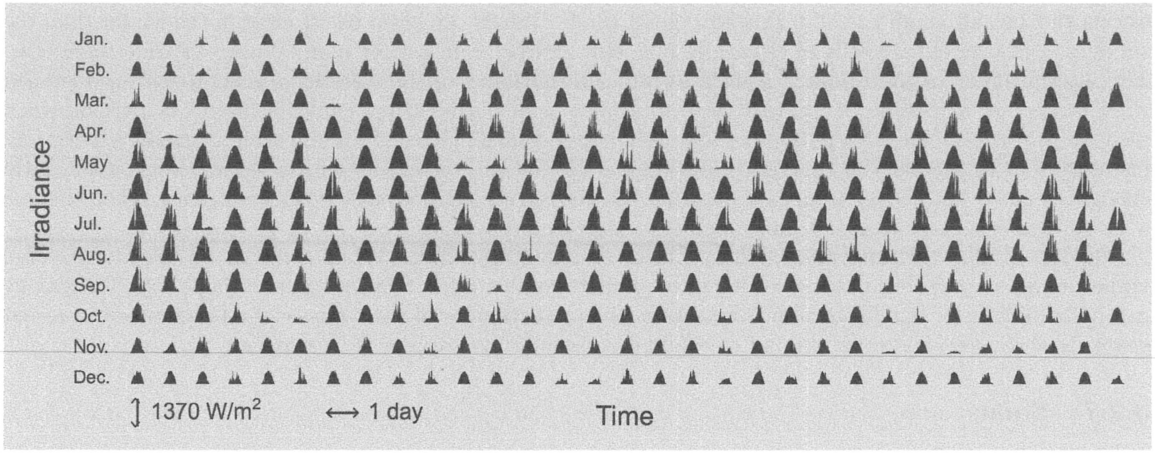

Figure 3.3: Complete Daily Solar Irradiance Profile Throughout the Full Year 2012 in Golden, Colorado.

Note: The horizontal timescale includes nights. Image credit: MIT Future of Solar Energy study, 2015, figure A.6. Reprinted with permission.

residential rooftops. Other factors may also reduce efficiency at smaller installations. For example, modules mounted on flat ground or a flat commercial rooftop can more easily be oriented for maximum exposure to the sun, than modules placed on a sloped residential rooftop.

The potential generation of a PV installation is determined by the installed capacity and by the amount of the solar resource, the **solar irradiance**. As with wind, the intermittency of the resource is a key challenge. The amount of solar irradiance varies dramatically. This happens on different timescales. The most obvious are the daily and the annual cycles. This also happens at much shorter timescales. "Solar irradiance can drop by a factor of five or more in the span of minutes as a result of passing clouds. The difference between a completely sunny day and a completely overcast day can amount to a 15-fold difference in integrated irradiance, and weather systems that produce overcast conditions sometimes persist for several days."[11] Figure 3.3 shows these different timescales of intermittency throughout a calendar year at a single location in the U.S. state of Colorado. The daily pattern shows clearly in each entry. The calendar effect of shorter days in the winter can be seen by contrasting the rows. The shorter timescale impact of clouds shows up episodically with spikes up or down in irradiance on a given day and in sequences of days that have markedly lower irradiance than the surrounding days. The intermittency of the resource translates into intermittency of the electricity generated from PV installations. Curtailment of generation is possible, but not often done.

The daily and annual cycles are predictable, but the episodic impact of cloud cover is less predictable. Some of the dramatic impact of cloud cover at a single location occurs on the timescale of minutes and can be averaged out over a larger geographic territory with dispersed solar installations. However, some of the impact of cloud cover is due to weather systems that cover larger territories and will not average out. These weather systems may be predictable at some time horizons, but not at others.

Solar Thermal or Concentrating Solar Power

Another form of solar power is **solar thermal**, or **concentrating solar power** (CSP). A CSP plant is a large-scale installation, usually with a capacity of 100 MW or more, that concentrates solar radiation, captures the heat and stores it in a fluid, and then uses this heat to produce steam and drive a turbine and generator. CSP plants require strong direct solar radiation, free from the scattering effect of clouds, haze and dust. This dramatically limits the siting possibilities. Because of the low density of energy in sunlight, CSP requires a large volume of land. Finally, CSP requires a high volume of water for cooling, which is often in limited supplies where siting for direct solar radiation is most desirable. An advantage of CSP is that the heat can be stored for a period of time, so that the time profile of electricity generation can be adapted to when the electricity is more needed, such as peak evening hours following sunset. Also, since the power is produced in a typical steam and turbine system, a CSP system can be integrated into a hybrid configuration with, for example, a natural gas-fired system in a combined-cycle plant which achieves higher thermal efficiency. The technology is still in its infancy and only a small amount of capacity is operational.

---

[11] MIT Future of Solar, p. 259.

3.3 Performance Metrics

3.3.1 Thermal Efficiency and Heat Rate

An important feature of many generation technologies is the **thermal efficiency** of the unit. This is the ratio of the energy output to the energy input:

$$\text{thermal efficiency} = \frac{\text{energy output}}{\text{energy input}}. \tag{3.1}$$

The theoretical maximum value for thermal efficiency is 100% since the energy output cannot be more than the energy input. In practical terms, rates of thermal efficiency for newly installed power plants typically range from 20% to 65% depending upon the technology.

Higher efficiency saves on fuel cost and reduces the quantity of pollutants produced per unit of electricity generated. Recent advances in materials and fabrications have made it possible to design PC units to operate at higher steam temperature and pressure, which raises their efficiency. This is what is described by the movement from subcritical to supercritical to ultra-supercritical steam parameters. New subcritical PC plants have thermal efficiencies in the range of 33-37%. Supercritical PC plants have efficiencies of 37-40%, and ultra-supercritical plants have efficiencies of 44-46%. Gas-fired combustion turbines achieve energy conversion efficiencies ranging between 20-35%.[12] Natural gas combined-cycle units usually operate with thermal efficiencies of between 50-60%.[13]. Recent improvements in turbine materials and technology have enabled efficiency improvements up to 65%.

For many purposes, it is useful to invert the efficiency ratio and report the amount of energy input required to produce a unit of electricity output, which is the **heat rate**:

$$\text{heat rate} = \frac{\text{energy input}}{\text{electricity output}} \frac{(\text{fuel units})}{(\text{electricity units})}. \tag{3.2}$$

However, in calculating a heat rate, we must take care about the units employed. While thermal efficiency is always reported using the same units in the denominator and numerator–such as the MWh of electricity output and the MWh of heat input–so that the reported value is unitless, heat rates are generally reported using different units for the energy input in the numerator and the energy output in the denominator. To emphasize this difference, equation 3.1 characterizes both the numerator and denominator as energy and makes no mention of units, while equation 3.2 explicitly distinguishes the numerator as energy and the denominator as electricity and mentions the different units used in the numerator and denominator. In much of the world where SI units are used, a heat rate is typically reported with kilojoules (kJ) in the numerator and kWh in the denominator. In some places, where the old metric system is still used, a heat rate is reported with kilocalories (kcal) in the numerator and kWh in the denominator. In the U.S. and other places where the British Imperial system of units is still used, a heat rate is typically reported with Btu in the numerator and kWh in the denominator. Therefore, to translate from thermal efficiency to heat rate, we

---

[12] U.S. Department of Energy, How Gas Turbine Power Plants Work, http://energy.gov/fe/how-gas-turbine-power-plants-work, accessed June 2016

[13] http://electrical-engineering-portal.com/an-overview-of-combined-cycle-power-plant

Table 3.1: The Heat Rate of a Power Plant with a 40% Efficiency Expressed in Four Different Units

Unit for Energy Input (numerator)		conversion factor	Heat Rate	
kilowatt hours	kWh	1.00	2.500	kWh/kWh
kilojoules	kJ	3,600.00	9,000	kJ/kWh
kilocalories	kcal	860.42	2,151	kcal/kWh
British thermal units	Btu	3,412.14	8,530	Btu/kWh

not only invert the ratio, but also multiply by the conversation factor appropriate to the units used for the energy input:

$$\text{heat rate} = \frac{1}{\text{thermal efficiency}} \cdot \text{units conversion factor}. \tag{3.3}$$

Table 3.1 shows the heat rate corresponding to a 40% thermal efficiency where the energy input is quoted in four different units: The second row shows the heat rate with the energy input measured in kilojoules: 1 over 40% equals 2.5, then multiplied by 3,600 kJ/kWh equals a heat rate of 9,000 kJ/kWh.

Since the heat rate is the inverse of the thermal efficiency, an improvement in the thermal efficiency corresponds to a reduction in the heat rate. This is shown in Figure 3.4. A lower heat rate means less fuel is required to produce the same unit of electricity. A higher heat rate means more fuel is required.

Table 3.2 shows the average heat rates for coal-, oil- (petroleum), natural gas-, and nuclear-fired power plants in the U.S. in recent years. Note the declining average heat rates for natural gas plants, indicating their improved efficiency over time.

Of course, the heat rate of individual plants within each category can vary significantly from the average. The heat rate varies depending upon a variety of factors, including steam temperature and pressure, and condenser cooling water temperature. For example, a unit in Florida will generally have a lower operating efficiency than a unit in northern New England or in northern Europe due to the higher cooling water

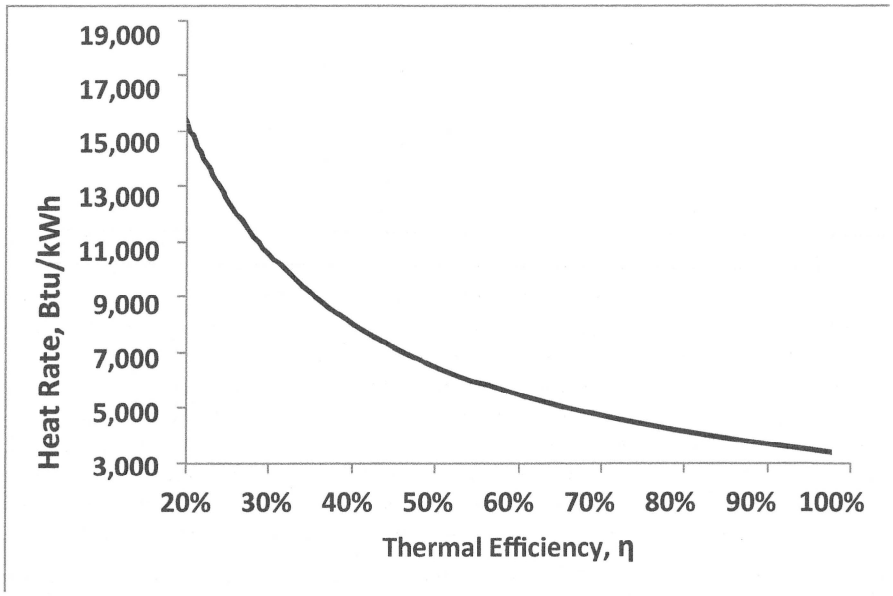

Figure 3.4: The Relationship Between Thermal Efficiency and Heat Rate

Table 3.2: Average Operating Heat Rate for Selected Energy Sources (Btu/kWh).

Year	Coal	Petroleum	Natural Gas	Nuclear
2004	10,331	10,571	8,647	10,428
2005	10,373	10,631	8,551	10,436
2006	10,351	10,809	8,471	10,435
2007	10,375	10,794	8,403	10,489
2008	10,378	11,015	8,305	10,452
2009	10,414	10,923	8,160	10,459
2010	10,415	10,984	8,185	10,452
2011	10,444	10,829	8,152	10,464
2012	10,498	10,991	8,039	10,479
2013	10,459	10,713	7,948	10,449
2014	10,428	10,814	7,907	10,459

Data source: U.S. Energy Information Administration.

temperature in Florida. The difference in generating efficiency could be 2 to 3 percentage points. Efficiency is higher in winter and at lower altitudes. For coal units, in particular, the heat rate depends upon the energy and carbon content of the coal, as well as upon the moisture, ash and sulfur content. While boilers can be optimized for a specific coal, as in the case of mine-mouth generating plants, most boilers are designed to accept a range of coals. Coals with lower energy content and high moisture can significantly reduce efficiency. However, they often come with a lower price which can make up for the lower efficiency. High sulfur content reduces efficiency due to the added energy consumption and operating costs required to remove the SO_2. Older plants are generally less efficient (have higher heat rates), both because they were built before efficiency-boosting technological improvements were introduced, but also simply because of the age of their equipment. Heat rates also vary depending upon how the unit is operated, as we discuss below. Figure 3.5 shows a histogram of the reported heat rates for all of the coal-fired power plants operating in the U.S. in 2012. They range from below 9,000 Btu/kWh to above 14,000.

3.3.2 Operating Cost and Marginal Operating Cost

The management and operation of generation plants must take careful consideration of cost. There are many different types of managerial decisions, each of which must take account of the right set of costs. When making investment decisions, it is important to evaluate all future costs, including both the capital expenditures required to install the capacity and the forecasted operating costs. An assessment that includes both up front capital costs and operating costs is a **total cost** calculation. It is also sometimes called a **long-run cost** calculation. Once the capacity has been built and the investment expenditures have been sunk, it is important to operate the unit taking into account current operating costs while disregarding sunk costs like the historical cost of the capacity. As assessment that includes only the operating costs, but only them, is an **operating cost** calculation. It is also sometimes called a **short-run cost** calculation.

In many contexts, where the relevant decision problem is clear, people refer simply to cost, without making clear whether they are looking at total or operating cost. So long as one understands that the guiding principle is the relevance of the cost item to the decision at hand, it should be straightfoward to resolve any ambiguity or disagreement. For example, while a short-run cost calculation does not include allocations of the capital costs originally paid to build the plant, it may nevertheless include charges for the incremental wear and tear on the equipment due to the actual generation. By incremental, we mean the wear and tear that would not occur if the plant were idle this hour. The cost of incremental wear and tear are difficult to estimate reliably, but they are real. This is reflected in the fact that turbine warranties often explicitly specify the number of operating hours for which they are valid, and maintenance manuals sometimes specify the number hours of operation between inspections or replacement of certain parts or other

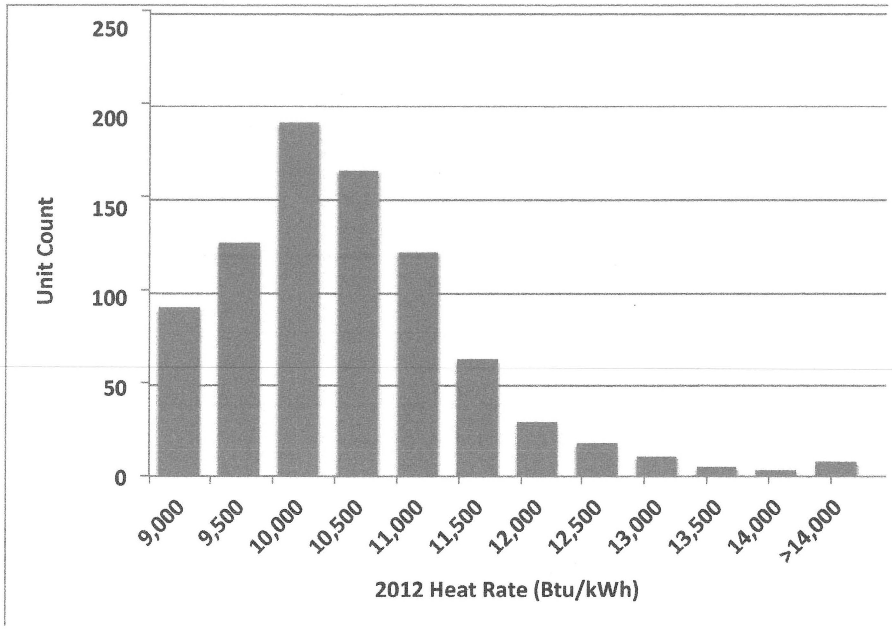

Figure 3.5: Reported Heat Rates at Coal-fired Power Plants in the U.S. in 2012.

Data source: U.S. Energy Information Administration, Analysis of Heat Rate Improvement Potential at Coal-Fired Power Plants, May 2015, Figure 4-6.

required procedures. In our discussion of steam turbine power plants above, we mentioned certain operating restrictions such as the maximum ramp rate and minimum downtime. While these are often discussed as fixed limits imposed by the technology, the units actually have some flexibility to operate outside the limits. A rational calculation of short-run marginal cost, including the wear and tear on the equipment, leads the operator to respect these limits in most cases, and to exceed the limits only in rare cases.

Figure 3.6 shows an illustrative operating cost curve for a steam turbine power plant with a maximum capacity of 500 MW. The plant also has a minimum generation of 100 MW. The formula for the cost curve is:

$$OC(q) \;=\; A \;+\; B \cdot q \;+\; C \cdot q^2 \tag{3.4}$$

where the variable q represents the level of generation this hour measured in MWh, and parameterized with A=\$4,365, B=\$34.11, and C= \$0.10. The curve shown in Figure 3.6 reflects only how cost varies with this hour's operation, taking as given all of the other decisions leading up to this hour. It is a short-run cost curve.

The **marginal cost** of electricity is the increment to cost resulting from an increment to generation: how much does total cost increase when one extra unit of generation is produced.

$$\text{marginal cost} \;=\; \frac{\Delta\ \text{Cost}}{\Delta \text{q}} \tag{3.5}$$

The concept of a marginal cost can be applied to any version of cost, whether short-run or long-run cost and whether just operating cost or total cost inclusive of capital cost. Since we have been discussing operating costs, it makes sense to consider the **marginal operating cost**. Starting from the operating cost curve defined in equation 3.4, the corresponding marginal operating cost equation is:

$$MOC(q) \;=\; \$34.11 \;+\; \$0.20 \cdot q \tag{3.6}$$

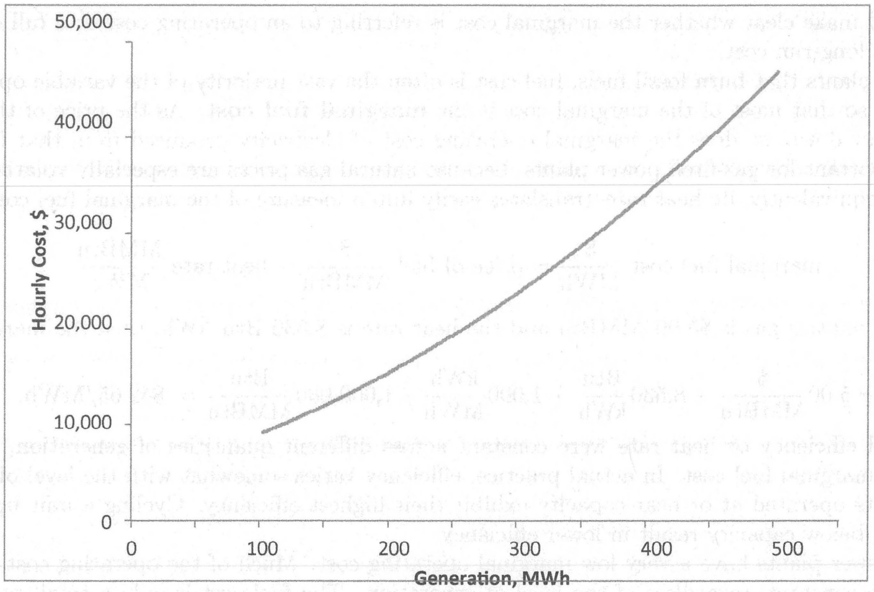

Figure 3.6: Illustrative Operating Cost Curve for a 500 MW Steam Turbine Power Plant.

Figure 3.7 shows this marginal operating cost as a thick gray line. In this case, the marginal operating cost is increasing slightly with the amount of generation–each additional unit adds a bit more to cost than the last unit added. It is common to use the term marginal cost without any modifier like 'operating'. The

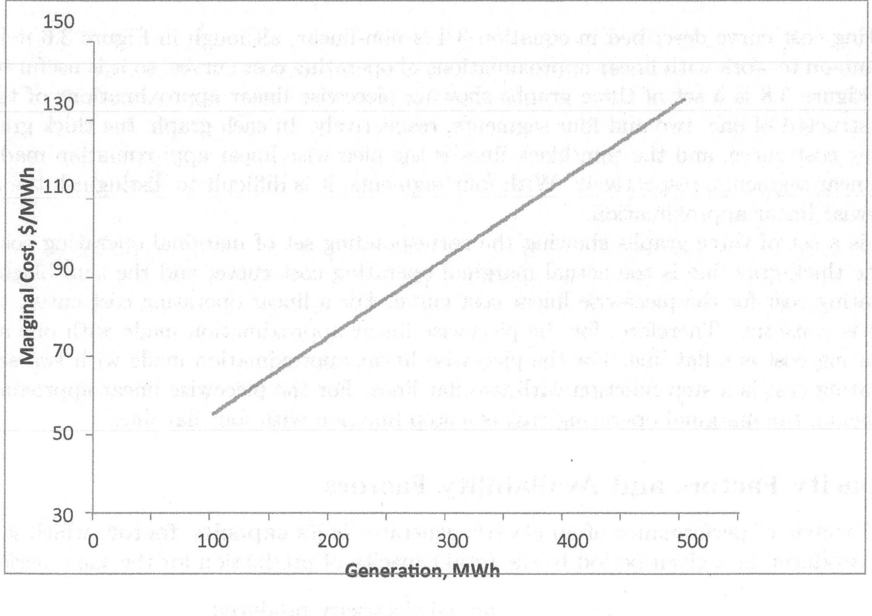

Figure 3.7: Illustrative Marginal Operating Cost Curve for a 500 MW Steam Turbine Power Plant.

context should make clear whether the marginal cost is referring to an operating cost or a full cost, or to a short-run or a long-run cost.

For power plants that burn fossil fuels, fuel cost is often the vast majority of the variable operating cost of generation, so that most of the marginal cost is the **marginal fuel cost**. As the price of the fossil fuel fluctuates up or down, so does the marginal operating cost of electricity produced from that fuel. This is especially important for gas-fired power plants, because natural gas prices are especially volatile. A plant's efficiency–or, equivalently, its heat rate–translates easily into a measure of the marginal fuel cost:

$$\text{marginal fuel cost } \frac{\$}{\text{MWh}} = \text{price of fuel } \frac{\$}{\text{MMBtu}} \cdot \text{heat rate } \frac{\text{MMBtu}}{\text{MWh}} \tag{3.7}$$

If the price of natural gas is \$5.00/MMBtu and the heat rate is 8,530 Btu/kWh, then the marginal cost of fuel is:

$$= 5.00\frac{\$}{\text{MMBtu}} \cdot 8{,}530\frac{\text{Btu}}{\text{kWh}} \cdot 1{,}000\frac{\text{kWh}}{\text{MWh}} \;/\; 1{,}000{,}000\frac{\text{Btu}}{\text{MMBtu}} = \$42.65/\text{MWh}. \tag{3.8}$$

If the thermal efficiency or heat rate were constant across different quantities of generation, then so too would be the marginal fuel cost. In actual practice, efficiency varies somewhat with the level of generation. Typically, units operated at or near capacity exhibit their highest efficiency. Cycling a unit up and down, and operating below capacity result in lower efficiency.

Nuclear power plants have a very low marginal operating cost. Much of the operating cost of a nuclear power plant is constant, regardless of the level of generation. The fuel cost is only a small portion of the operating cost. Wind and solar plants have a marginal cost that is approximately zero. The fuel is costless, and most of the other operating costs are constant, regardless of the level of generation. It is tempting to think of hydro plants as also having a costless fuel, and this may be a good rough approximation for a run-of-the-river plant. However, for a plant with a large reservoir and stock of water, the key operating decision will be whether to generate the power today or tomorrow. A unit of water used today cannot be used again tomorrow. So, at any specific point in time, there is an opportunity cost to using the water today. Even for a run-of-the-river plant, the diversion of water through the turbines is a choice of how to use the water, and can sometimes involve tradeoffs. The history of not pricing water use can obscure the real nature of these costs.

The operating cost curve described in equation 3.4 is non-linear, although in Figure 3.6 it looks almost linear. It is common to work with linear approximations of operating cost curves, so it is useful to appreciate their nature. Figure 3.8 is a set of three graphs showing piecewise linear approximations of the operating cost curve constructed of one, two and four segments, respectively. In each graph, the thick gray line is the actual operating cost curve, and the thin black lines is the piecewise linear approximation made up of one, two and four linear segments, respectively. With four segments, it is difficult to distinguish the actual curve from the piecewise linear approximation.

Figure 3.9 is a set of three graphs showing the corresponding set of marginal operating cost curves. In each graph, the thick gray line is the actual marginal operating cost curve, and the thin black lines is the marginal operating cost for the piecewise linear cost curve. For a linear operating cost curve, the marginal operating cost is constant. Therefore, for the piecewise linear approximation made with one segment, the marginal operating cost is a flat line. For the piecewise linear approximation made with two segments, the marginal operating cost is a step function with two flat lines. For the piecewise linear approximation made with four segments, the marginal operating cost is a step function with four flat lines.

3.3.3 Capacity Factors and Availability Factors

A widely cited metric of performance of an electric generator is its **capacity factor**, which is the ratio of the electricity produced in a given period to the total capacity of production for the same period:

$$\text{capacity factor} = \frac{\text{actual electricity produced}}{\text{potential electricity production}}. \tag{3.9}$$

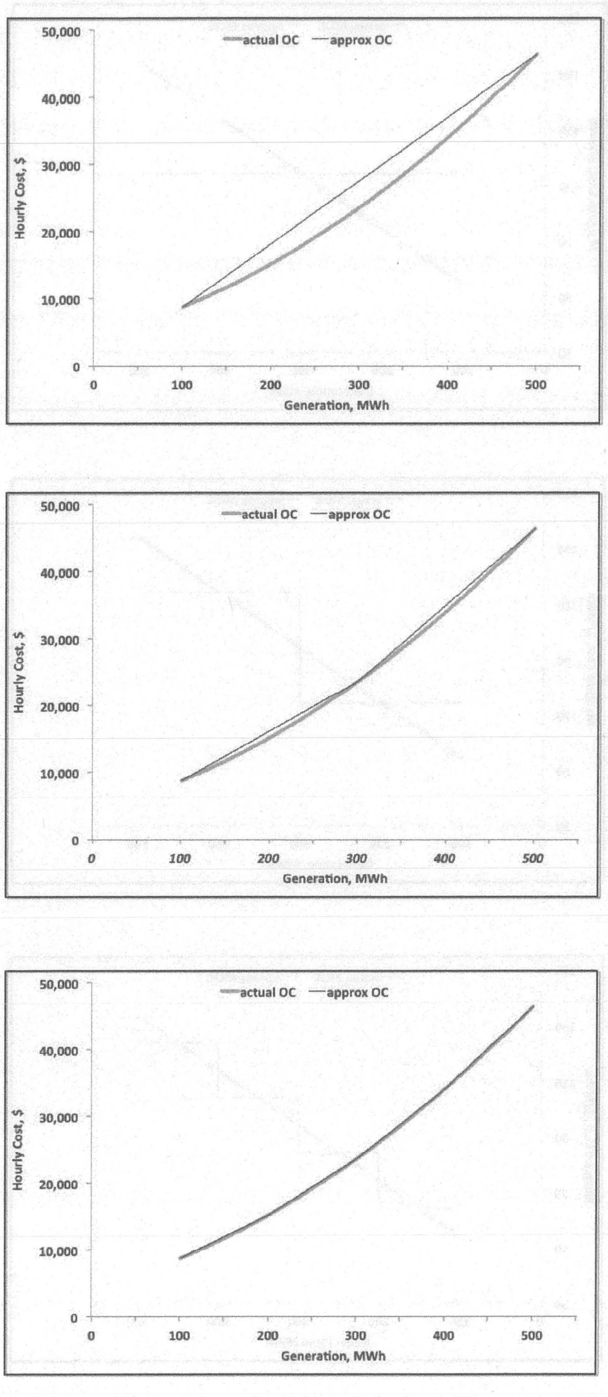

Figure 3.8: Piecewise Linear Approximations of an Operating Cost Curve With 1, 2 and 4 Segments.

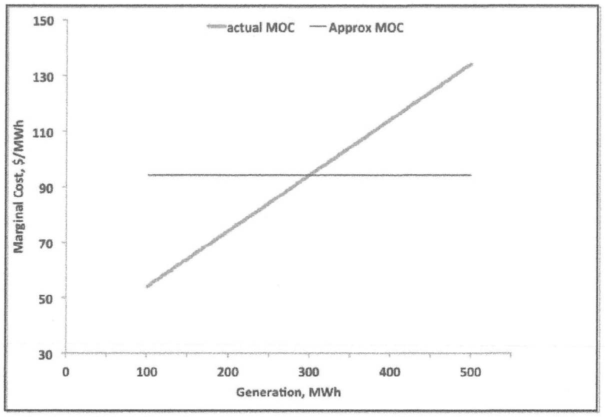

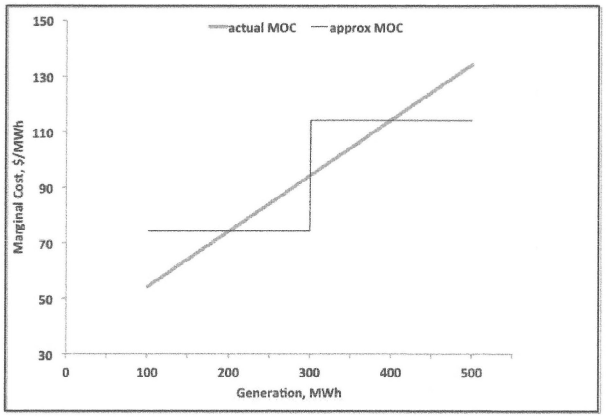

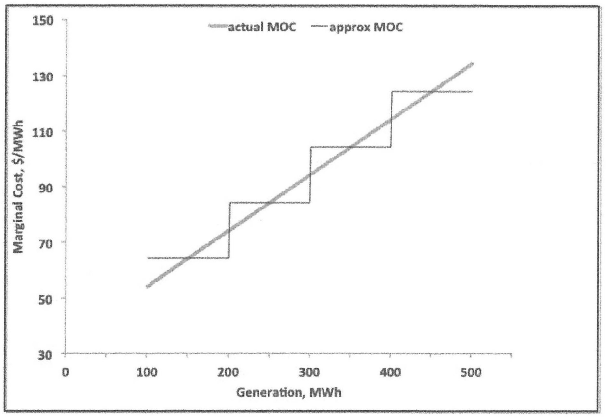

Figure 3.9: The Marginal Operating Cost Curves for the Piecewise Linear Approximations With 1, 2 and 4 Segments.

Table 3.3: Illustrative Calculation of Availability Factor and Capacity Factor for a One-Day Period.

	Capacity		
Hour	Nameplate	Available	Generation
1	500	500	0
2	500	500	100
3	500	500	200
4	500	500	300
5	500	500	400
6	500	500	500
7	500	500	500
8	500	500	500
9	500	500	500
10	500	500	500
11	500	500	500
12	500	500	500
13	500	500	500
14	500	500	500
15	500	500	500
16	500	500	500
17	500	500	500
18	500	500	400
19	500	500	300
20	500	500	200
21	500	500	100
22	500	0	0
23	500	0	0
24	500	0	0
Total	12,000	10,500	8,000
Availability Factor=		87.5%	
Capacity Factor=			66.7%

Table 3.4: Average Annual and Monthly Capacity Factors at Coal-fired Generators in the U.S..

Annual Factors	
2013	59.7%
2014	61.0%
2015	54.7%
Year 2015	
January	61.3%
February	64.9%
March	50.3%
April	43.3%
May	49.8%
June	62.6%
July	66.8%
August	64.9%
September	58.7%
October	47.0%
November	43.9%
December	43.6%

Data source: U.S. Energy Information Administration.

A capacity factor is a unitless ratio, typically reported as a percentage. Table 3.3 illustrates the calculation of the capacity factor for a power plant with a capacity of 500 MW over a 24-hour period. If the plant operated at full capacity for all 24 hours, its production would be 12,000 MWh, so this is the potential electricity production. However, the plant is not on at all in the first hour, and then ramps up in the early morning, producing less than its full capacity for several hours. In the late evening, the plant ramps down and is shut-off. The actual electricity production is only 8,000 MWh. The capacity factor is the ratio of 8,000/12,000= 66.7%.

A generator typically has what is called its **nameplate capacity** or rated capacity, which reflects what the manufacturer certifies the unit can produce on the date of installation under some set of standard conditions, including the ambient temperature and the local air pressure. An individual generator's actual capacity depends upon how conditions at its location vary from the standard. For example, the Genkai 4 nuclear plant in Japan has occasionally reported a capacity factor just above 100% because the ambient conditions at its location were better than the standard conditions used to set the nameplate capacity. Many generators have separate capacity ratings for the summer and the winter to capture the impact of the standard temperature variation on the capacity. The capacity factor can be calculated gross or net of power used at the generating station. In different regions and countries, standard setting agencies specify the precise terminology and methodology for calculating the capacity factors relied upon by system operators and other parties.

Table 3.4 shows some average annual and monthly capacity factors for coal-fired generators in the U.S., which range between 43 and 67%.

A plant's actual capacity can decline over its years of operation, or it can be boosted through investment. For certain purposes, it is useful to re-designate the unit's capacity based on historical performance or other considerations. This is especially important when evaluating whether the entire system has adequate capacity–as we discuss in Chapter 5–and when running a capacity market–as we discuss in Chapter 8.

Another widely cited metric is a unit's **availability factor**, which is the ratio of the amount of time when a generating unit is available to produce electricity to the total amount of time in the period:

$$\text{availability factor} = \frac{\text{time available to produce}}{\text{total time}}. \tag{3.10}$$

The availability factor is a unitless ratio, typically reported as a percentage. It must be less than or equal

Table 3.5: Average Annual Capacity Factors for Utility Scale Non-Fossil Fueled Generators in the U.S.

Year	Nuclear	Conventional Hydropower	Wind	Solar PV	Landfill Gas
2013	89.9%	38.9%	32.4%	NA	68.9%
2014	91.7%	37.3%	34.0%	25.9%	68.9%
2015	92.3%	35.8%	32.2%	25.8%	68.7%

Data source: U.S. Energy Information Administration.

to one. The hours that a generator is off-line for maintenance reduce the availability factor. Under certain circumstances, a generating unit may be available to produce some but not all of its capacity. This can be included in the calculation of an equivalent availability factor: for example, an hour when the generator is only available to produce at 1/2 capacity is included as a 1/2 hour. In different regions and countries, standard setting agencies specify the precise terminology and methodology for calculating the availability factors relied upon by system operators and other parties. Table 3.3 illustrates the calculation of an availability factor and how this contrasts with the capacity factor. In the late evening, when the plant ramps down and is shut-off, it must remain off for some hours. Therefore, its capacity is no longer available. The available electricity capacity for the entire day is only 10,500 MWh. The availability factor is the ratio of 10,500/12,000= 87.5%.

There are many other additional metrics, each of which captures a different aspect of a plant's operation. For example, it may be useful to distinguish between periods when a generating unit is unavailable due to planned maintenance and periods when a generating unit is unexpectedly unavailable due to a **forced outage**, which is an unplanned, unexpected shutdown. Therefore, some systems require generators to report a forced outage rate.

Table 3.5 shows recent average annual capacity factors for non-fossil fueled generators in the U.S., including nuclear and different types of renewable facilities. The average capacity factor of nuclear facilities is much higher than the average capacity factors at coal facilities shown in Table 3.4. On the other hand, the average capacity factor of hydropower, wind and solar facilities is much lower. This is primarily due to the intermittency of the renewable resource. Calculations of the capacity assume that the unit is producing at all hours of the day, regardless of the flow of water or wind or of the solar insolation. For example, the simple fact that the sun is out only part of the day means that actual insolation is dramatically less than the potential value used in the denominator of a capacity factor calculation. Even a prime location in southern California might receive approximately 5.8 kilowatt-hours of solar radiation per square meter per day ($\text{kWh/m}^2\text{/day}$), which is already just 23.9% of the potential value.

The capacity factor at individual hydro plants vary widely around this average. A plant's capacity factor will be shaped by the seasonal pattern of water flow, and by what sort of electricity services the facility provides. One of the main advantages of dam systems is the ability to adjust power generation very easily. This enables a hydropower plant to provide different kinds of electricity services, such as frequency regulation, and to act as a reserve. As we shall see in later chapters, these may be valuable services although they entail a lower volume of actual electricity generation relative to the plant's capacity. The different seasonal flow patterns of the two rivers where China's Three Gorges plant and Brazil's Itaipu plant are located mean that the Three Gorges' capacity factor is only 44% as compared to the Itaipu's 75%. Indeed, Itaipu's total generation of 89,500 GWh slightly exceeded the Three Gorges' generation of 87,000 GWh, despite the fact that Itaipu's capacity is so much less than the Three Gorges'.[14]

Wind farm capacity factors, which average around 35%, can range from just below 10% to just above 50%.

As with other generation technologies, the potential capacity of a solar PV installation is calculated based on standard test conditions. Actual conditions vary dramatically with geography. A main determinant is

---

[14] U.S. Energy Information Administration, The world's nine largest operating power plants are hydroelectric facilities, Today in Energy, October 18, 2016.

Table 3.6: Various Services Provided by Storage

Generation Service	Time Scale
regulation	Seconds
operating reserves	Minutes
black start	
peak shaving	Minutes/Hour
energy arbitrage / energy firming	Hours/Days
Transmission & Distribution Service	
back-up power	
peak shaving	Minutes/Hour

latitude on the globe, but other factors contribute as well. Chief among them is cloud cover, as well as the moisture content of the air and various pollutants. For example, a typical location in the northeastern state of Massachusetts might receive approximately 3.8 kWh/m^2/day of insolation, which is 1/3 less than the 5.8 kWh/m^2/day of a prime southern California location.

3.4 Storage

A critical difference between electricity and other commodities is that electricity cannot be stored. An electrical system must be carefully operated so that total generation almost exactly matches total demand at every instant in time. The system is extremely unstable with respect to any errors in matching the two. Lasting errors of very small magnitude can cause parts of the system to briefly go out of service, or, worse, cause large regions to collapse in a lasting blackout.

However, while electricity itself cannot be stored as such, electricity can be converted into various forms of stored energy which can be quickly converted back into electricity. In a pumped storage system, electricity is used to pump water uphill where the water can sit in a reservoir until it is released and used to power a turbine and generator, recreating electricity. In a battery system, electricity is used to create the chemical charge that is later discharged to recreate electricity. These are all forms of storage similar to what exists for other commodities where the commodity itself is stockpiled, whether grain in a silo, oil in a tank, or natural gas in an underground cavern. Consequently, although the electricity itself is never actually stockpiled, we commonly speak of these systems as electricity storage.

At present, electricity storage plays a very small role in electricity systems. The U.S., for example, has less than 25 GW of capacity from storage of all kinds. That compares to a total generation capacity that exceeds 1,000 GW, so the total capacity from storage equals approximately 2% of the country's total generation capacity. This storage serves many different functions, as indicated in Table 3.6. The functions are often distinguished by the time scale over which the service is provided. Regulation, which is an ancillary service discussed in Chapter 4, requires very, very quick and short bursts of energy at a time scale of seconds. Operating reserves, another ancillary service discussed in Chapter 4, require a response in minutes, and one that can last for a good number of minutes. **Peak shaving** involves supplying energy at those moments in the day when the system is straining to provide enough generation. If storage can shift even a small amount of power just a small amount of time, it may help prevent a loss of load event, or it may substitute for expensive investments in generating capacity that would otherwise be needed to prevent the event. When most people think of storage, they often think of longer term **energy arbitrage**: buying power in hours when demand is low and power is cheap, and selling power when load is high and power is expensive. This includes buying power at night and selling it during the day. This is also called **energy firming** since the large price differentials it targets are often driven by the large capacity investments in intermittent renewable energy resources. Table 3.6 also provides examples of transmission and distribution services which benefit from storage. Storage can provide small amounts of back-up power for key services at a few locations on a network and thereby provide a bridge while other actions are taken. In addition, storage can temporarily substitute for investments in network reinforcements by reducing peak load. Different storage technologies are suitable for different services, as we explain below. Some technologies can respond very quickly, but cannot store high volumes of power. These are suitable for regulation, but not for energy arbitrage. Only a few technologies are currently suitable for arbitrage or firming. Also keep in mind that the figure cited above for total storage capacity is small compared to the predictable daily and seasonal variation in load and also compared to the unpredictable variation we saw in Chapter 2. Storage currently plays a small role because it remains a very expensive option. For now, building reserve capacity in generators and stockpiling fuel remain competitive alternatives for coping with the predictable and unpredictable variation in demand. And, despite the cost of reserve capacity, it is currently often the cheapest alternative.

Storage may play a larger role in the electricity system of the future. Scientists are working to lower the costs and improve the performance of storage systems. At the same time, the increasing necessity of reducing CO_2 emissions and the increasing use of intermittent renewables may increase the demand for storage. Already authorities in some regions are requiring the installation of larger volumes of storage. How significant this change will be is something that only time can tell.

3.4.1 Pumped Hydro

A pumped hydro facility has two reservoirs–one at a lower elevation and one at a higher elevation. In times of low electricity demand, electricity is used to pump water from the lower reservoir up to the upper reservoir. In times of peak electricity demand, the water flows from the upper reservoir down to a turbine connected to an electrical generator.

The rate at which a storage facility can generate electricity is often called its **power capacity**, and the aggregate amount of electricity that can be delivered before the storage is exhausted is called its **energy capacity**. In a pumped hydro facility, the scale of the turbine and generator determine the maximum power capacity, and the scale of the reservoirs determine the maximum energy capacity. The **discharge time** is the amount of time the system can deliver electricity. If the power capacity and energy capacity were constant as the system is drained, then the discharge time would be the ratio of the energy capacity to the power capacity. Pumped hydro facilities can be very large, with generation capacities in the hundreds of megawatts and even gigawatts, and their discharge time is sometimes over tens of hours before the reservoir is exhausted. The high power capacity and long discharge time make them especially useful for managing large fluctuations in load lasting over a timescale of hours.

Another important metric for storage technologies is the **conversion efficiency**, which is the ratio of electricity output to electricity input. A pumped storage facility can have a conversion efficiency of over 80%.[15] The precise level of the conversion efficiency depends upon how the system is operated. Water in the upper reservoir evaporates with time, so that a system that cycles the water daily, generating a high amount of electricity from the same reservoir, will have a higher conversion efficiency than one that cycles the water only occasionally, generating a smaller amount of electricity from the reservoir. The one that cycles infrequently will see a larger fraction of each unit of pumped water lost to evaporation.

Pumped storage represents the vast majority of grid-scale electricity storage. In the U.S., it represents about 98% of the storage capacity.

A major limitation on pumped storage facilities are siting issues: environmental considerations often limit construction of new facilities at locations that are otherwise ideal.

3.4.2 Compressed Air Systems

The first utility scale compressed air energy storage (CAES) system, the Huntorf plant, began operation near the port city of Bremen in Germany in 1978. In times of low electricity demand, electricity is used to drive a compressor that pumps air from the atmosphere into two caverns which have a total volume of approximately 300,000 m^3. It takes about 8 hours to fill the caverns. In times of peak electricity demand,

---

[15] International Energy Agency (IEA) & International Renewable Energy Agency (IRENA) Hydropower Technology Brief, 2015.

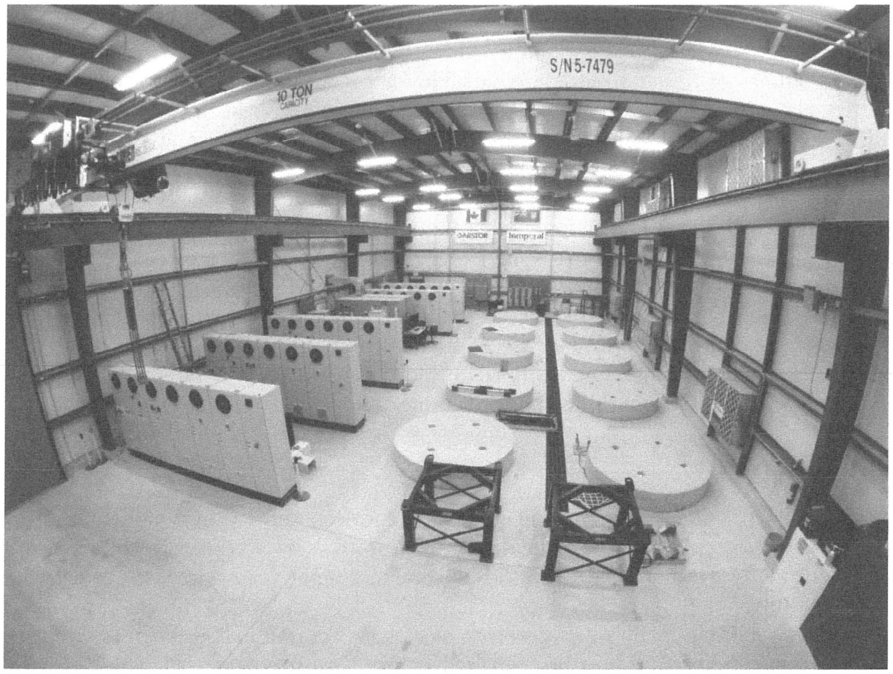

Figure 3.10: NRStor's Flywheel Facility in Minto, Ontario, Canada, with 2 MW of Temporal Power Flywheel Energy Storage.

Image credit: NRStor website. Reprinted with permission.

the air is released into a two-stage gas combustion turbine, providing precompressed air. This triples the useful power produced by the combustion turbine.[16]

Other designs have been and continue to be developed that utilize the compressed air without the use of a fossil fuel-fired generator.

The distinctive feature of compressed air systems is that, like pumped storage, they have both high power capacities and long discharge time, providing large volumes of power over several hours which are useful in managing the daily cycle of demand or large fluctuations in demand that last a number of hours. Research is ongoing to develop better utility-scale flow batteries that can compete with compressed air or pumped hydro systems, but they are not available now.

3.4.3 Flywheels

Flywheels are an age old technology for storing energy as rotational energy. Figure 3.10 shows a facility in Ontario, Canada housing flywheels with 2 MW of capacity. Flywheels are able to provide strong instantaneous bursts of power which are useful for regulating the minute-by-minute balancing of an electricity system, and they recharge very quickly. However, they have a short discharge time, and so cannot provide lasting injections of power. Older systems have demonstrated a conversion efficiency in the 30-40% range. Newer systems outfitted with equipment to reduce friction have achieved efficiencies as high as 90%. Flywheel systems have a very, very long lifetime and are very, very low maintenance.

---

[16] H. Hoffeins, "Huntorf Air Storage Gas Turbine Power Plant," Brown Boveri, Publication No. D GK 90 202 E, 1994.

3.4.4 Batteries

The battery was invented in 1799, and has been used on telegraph networks and the electric grid since the mid-1800s. Modern batteries have a discharge time from a few seconds up to 6 hours, depending upon the chemistry employed and other factors.[17] This discharge time lies between the long time frames of pumped hydro and CAES and the very short time frame of flywheels. They are generally used to provide support to the transmission and to the distribution system, helping assure stability and security. There is a growing demand for these services, and therefore new investments in battery deployment. Load can sometimes use a battery to shave their peak demand, which may reduce the transmission or distribution bill. They are not generally used to provide significant lasting reserve generating capacity, as CAES and pumped hydro are. However, battery technology is an area of intense research in the hope of improving and expanding capacities and dropping costs, and the technology is quickly evolving. Services that batteries have not historically performed, may be performed by the batteries of the future.

---

[17] Huff, Georgianne, et al. DOE/EPRI Electricity Storage Handbook in Collaboration with NRECA. No. SAND2015-1002. Sandia National Laboratories (SNL-NM), Albuquerque, NM (United States), 2015.

Study Review

Some key terms introduced here:

- primary energy
- steam turbine (ST)
- dispatchable
- minimum start time
- minimum generation level
- ramp rate
- minimum downtime
- ranks of coal
- combustion turbine (CT)
- combined-cycle gas turbine (CCGT)
- configurations
- fuel deliverability
- dual-fired plants
- reservoir
- run-of-river
- head height
- flow rate
- specific power
- intermittency
- curtail
- solar photovoltaics (PV)
- solar irradiance
- solar thermal, or concentrating solar power (CSP)
- thermal efficiency
- heat rate
- total cost
- long-run cost
- operating cost
- short-run cost

- marginal cost
- marginal operating cost
- marginal fuel cost
- capacity factor
- nameplate capacity
- availability factor
- forced outage
- peak shaving
- energy arbitrage
- energy firming
- power capacity of storage
- energy capacity of storage
- discharge time
- conversion efficiency